

(12) UK Patent Application (19) GB (11) 2 347 252 (13) A

(43) Date of A Publication 30.08.2000

(21) Application No 0003269.8

(22) Date of Filing 14.02.2000

(30) Priority Data

(31) 09256279

(32) 23.02.1999

(33) US

(71) Applicant(s)

Motorola Inc

(Incorporated in USA - Delaware)

**Corporate Offices, 1303 East Algonquin Road,
Schaumburg, Illinois 60196, United States of America**

(72) Inventor(s)

David Erik Chevalier

Henry L Kazecki

(74) Agent and/or Address for Service

Peter D Hudson

**Motorola Limited, European Intellectual Property
Department, Midpoint, Alencon Link, BASINGSTOKE,
Hampshire, RG21 7PL, United Kingdom**

(51) INT CL⁷

G10L 15/20 // G10L 15:12 15:14 101:04 101:10

(52) UK CL (Edition R)

G4R RRL R1F

U1S S2123 S2127 S2204 S2215

(56) Documents Cited

GB 2188764 A

(58) Field of Search

UK CL (Edition R) G4R REX RRL

INT CL⁷ G10L 15/00 15/08 15/14 15/20

Online:WPI, EPODOC, JAPIO

(54) Abstract Title

Voice recognition system

(57) A voice recognition system generates a variable rejection strictness as a function of at least one background noise level measured during training and noise signal measurements made during an input utterance made during recognition mode of operation. A word entrance penalty is assigned as a function of the variable rejection strictness.




FIG. 6

At least one drawing originally filed was informal and the print reproduced here is taken from a later filed formal copy.

This print takes account of replacement documents submitted after the date of filing to enable the application to comply with the formal requirements of the Patents Rules 1995




FIG. 1




FIG. 2




FIG. 3


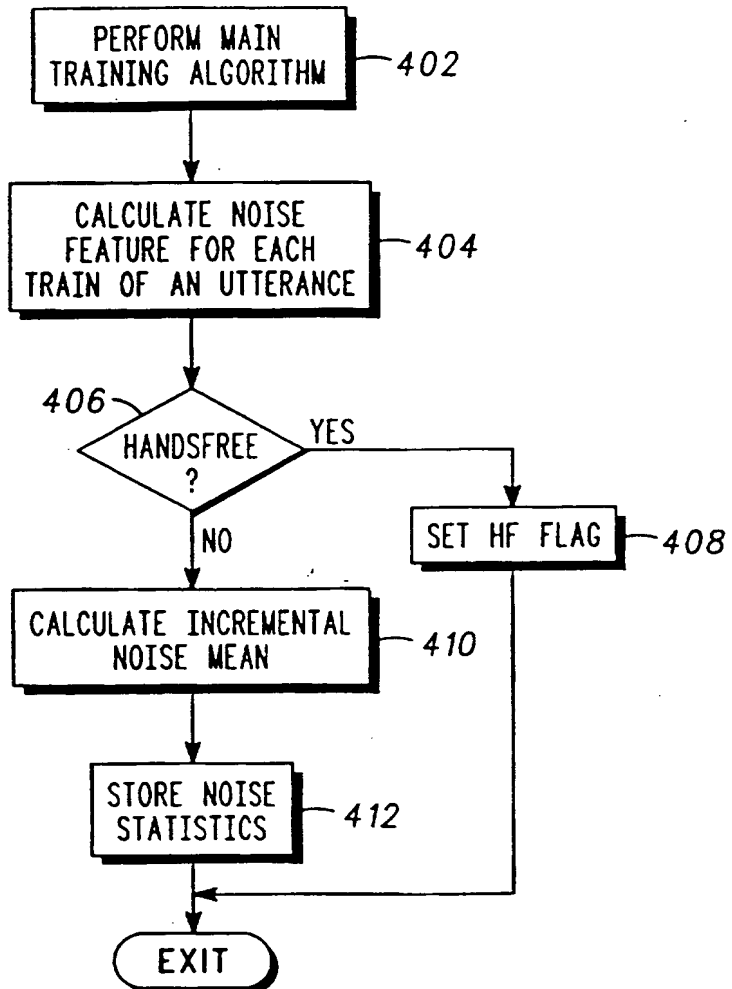


FIG. 8

*FIG. 4*




FIG. 6




FIG. 5

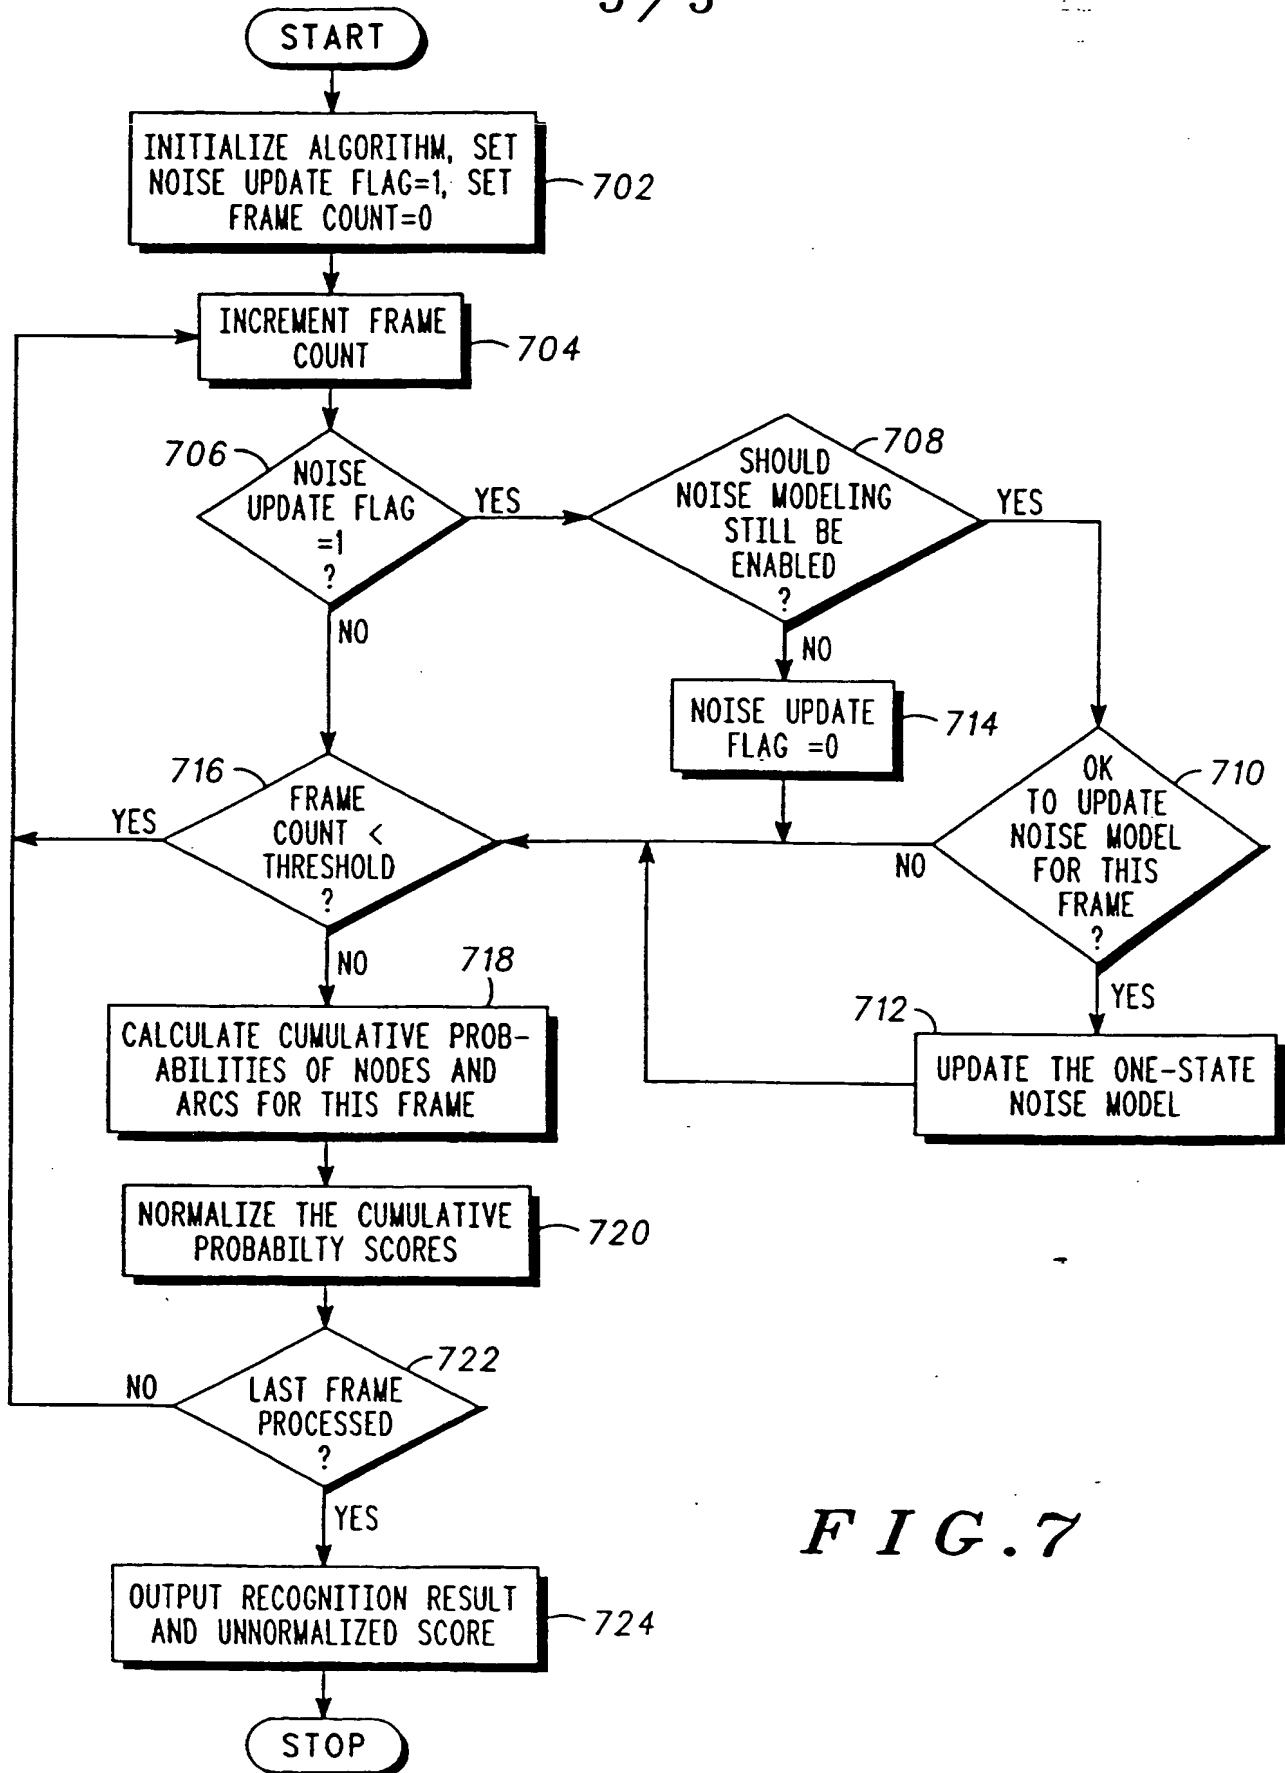


FIG. 7

VOICE RECOGNITION SYSTEM AND METHOD

FIELD OF THE INVENTION

5 The present invention pertains to voice recognition.

BACKGROUND OF THE INVENTION

 Speaker dependent speech recognition systems use a feature
10 extraction algorithm to perform signal processing on a frame of the
input speech and output feature vectors representing each frame. This
processing takes place at the frame rate. The frame rate is generally
between 10 and 30 ms, and will be exemplified herein as 20 ms in
duration. A large number of different features are known for use in
15 voice recognition systems.

 Generally speaking, a training algorithm uses the features
extracted from the sampled speech of one or more utterances of a word
or phrase to generate parameters for a model of that word or phrase.
This model is then stored in a model storage memory. These models
20 are later used during speech recognition. The recognition system
compares the features of an unknown utterance with stored model
parameters to determine the best match. The best matching model is
then output from the recognition system as the result.

 It is known to use a Hidden Markov Model (HMM) based recognition
25 system for this process. HMM recognition systems allocate frames of the
utterance to states of the HMM. The frame-to-state allocation that produces the
largest probability, or score, is selected as the best match.

 Many voice recognition systems do not distinguish between
valid and invalid utterances. Rather, these systems choose one of the
30 stored models which is the closest match. Some systems use an Out-of-
Vocabulary rejection algorithm which seeks to detect and reject invalid
utterances. This is a difficult problem in small vocabulary, speaker

dependent speech recognition systems due to the dynamic size and unknown composition of the vocabulary. These algorithms degrade under noisy conditions, such that the number of false rejections under noisy conditions increases.

5 In practice, out-of-vocabulary rejection algorithms must balance performance as measured by correct rejections of invalid utterances and false rejections of valid utterances. The false rejection rate can play a critical role in customer satisfaction, as frequent false rejections, like incorrect matches, will cause frustration. Thus, out-of-
10 vocabulary rejection is a balance of meeting user expectations for recognition.

 Accordingly it is known to calculate a rejection threshold based upon the noise level. For example, it is known to measure the noise level prior to the detection of the first speech frame. A threshold is
15 calculated from the measurement. An input is rejected if the difference between the word reference pattern and the input speech pattern is greater than the rejection threshold. Such a system is thus dependent upon an arbitrary noise input level. Such measurement can not be
20 relied upon to produce a meaningful rejection decision.

 Accordingly, there is a need for an improved method of providing a basis for rejecting utterances in a voice recognition system.

BRIEF DESCRIPTION OF THE DRAWINGS

25 FIG. 1 is a circuit schematic in block diagram form illustrating a wireless communication device.

 FIG. 2 is a circuit schematic in block diagram form illustrating a voice recognition system in the device according to FIG. 1.

 FIG. 3 is an illustration of a grammar network with two nodes.

30 FIG. 4 is a flow chart illustrating training.

 FIG. 5 illustrates a window and frames corresponding thereto.

 FIG. 6 is a high level flow chart illustrating recognition.

FIG. 7 is a flow chart illustrating training during recognition.
FIG. 8 illustrates the penalty function.

DETAILED DESCRIPTION OF THE DRAWINGS

5

The present invention has a variable rejection strictness depending upon the background noise levels during training and recognition. During training, noise features are generated from the training utterances. An incremental noise reference mean is updated
10 from the noise features. The statistics are stored in memory to make them available to the recognition algorithm. Noise statistics are not updated when training in a handsfree mode because of the higher levels of background noise. If there are no noise statistics available, the recognition algorithm defaults to the minimum strictness.

15 During recognition, the input noise energy feature is compared to the reference noise statistics and a noise ratio is computed. The strictness of the out of vocabulary rejection algorithm is then selected based upon the noise ratio. The present invention helps to prevent false rejection of valid utterances in the presence of noise.

20 The strictness parameter is a word entrance penalty in the two level alignment algorithm recognition search. The confidence measurement of the best path is implemented as a zero mean one state garbage model in parallel with the voice tag models.

A device 100, in which the invention can be advantageously employed
25 is disclosed in FIG. 1. The device 100 is described to be a portable radiotelephone herein for illustrative purposes, but could be a computer, a personal digital assistant, or any other device that can advantageously employ voice recognition, and in particular a device which can take advantage of a memory efficient voice recognition system. The illustrated radiotelephone
30 includes a transmitter 102 and a receiver 104 coupled to an antenna 106. The transmitter 102 and receiver 104 are coupled to a call processor 108, which performs call processing functions. The call processor 108 can be implemented

using a digital signal processor (DSP), a microprocessor, a microcontroller, a programmable logic unit, a combination of two or more of the above, or any other suitable digital circuitry.

5 The call processor 108 is coupled to a memory 110. Memory 110 contains RAM, electronically erasable programmable read only memory (EEPROM), read only memory (ROM), flash ROM, or the like, or a combination of two or more of these memory types. The memory 110 supports operation of the call processor 108, including the voice recognition operation, and must include an electronically alterable memory to support the state
10 transition path memory. The ROM can be provided to store the device operating programs.

 An audio circuit 112 provides digitized signals from a microphone 114 to call processor 108. The audio circuit 112 drives speaker 116 responsive to digital signals from the call processor 108.

15 The call processor 108 is coupled to a display processor 120. The display processor is optional if additional processor support is desired for the device 100. In particular, the display processor 120 provides display control signals to the display 126 and receives inputs from keys 124. The display processor 120 can be implemented using a microprocessor, a microcontroller, a
20 digital signal processor, a programmable logic unit, a combination thereof, or the like. A memory 122 is coupled to the display processor to support the digital logic therein. The memory 122 can be implemented using RAM, EEPROM, ROM, flash ROM, or the like, or a combination of two or more of these memory types.

25 With reference to FIG. 2, the audio signals received by microphone 114 are converted to digital signals in an analog-to-digital converter 202 of audio circuit 112. Those skilled in the art will recognize that the audio circuit 112 provides additional signal processing, such as filtering, which are not described herein for brevity. The call processor, 108, performs feature extraction 204 on
30 the processed digital signal representation of the analog signal output by microphone 114 and produces a set of feature vectors representative of the user utterance. A feature vector is produced for each short time analysis window.

The short time analysis window is a frame, which in the example illustrated herein is 20 ms. Thus there is one feature vector per frame. The processor 108 uses the features for speech recognition 206 or training 207.

In training, the feature vectors of the utterance are used to
5 create templates in the form of HMMs, which are stored in memory 208. In speech recognition, the feature vectors representing the input utterance are compared to the templates of stored vocabulary words in memory 208 to determine what the user said. The system may output the best match, a set of the best matches, or optionally, no match.
10 Memory 208, is preferably a non-volatile memory portion of memory 110 (FIG. 1), and may for example be an EEPROM or flash ROM. As used herein, "words" can be more than one word, such as "John Doe," or a single word such as "call".

The feature extractor 204, generally performs signal processing on a
15 frame of the input speech, and outputs feature vectors representing each frame at the frame rate. The frame rate is generally between 10 and 30 ms, and may for example be 20 ms in duration. Trainer 207 uses the features extracted from the sampled speech of one or more utterances of a word or phrase to generate parameters for a model of that word or phrase. This model is then stored in a
20 model storage non-volatile memory 208. The model size is directly dependent upon the feature vector length, such that a longer feature vector length requires a larger memory.

The models stored in memory 208 are then used during recognition
206. The recognition system performs a comparison between the features of an
25 unknown utterance and stored model parameters to determine the best match. The best matching model is then output from the recognition system as the result.

With reference now to FIG. 3, a grammar network representing speech
recognition is illustrated. The Node N_1 and N_2 are connected by HMM models
30 represented by arcs A_1 through A_N plus a garbage model arc A_{GM} . Arcs A_1 through A_N represent all of the individual HMM models that have been trained

in the voice recognition systems and stored in the memory 208. The garbage model arc represents a single state garbage model reference.

The node N_1 includes a single state noise model A_1^{noise} . The node N_2 similarly contains a single state noise model A_2^{noise} . The recognition system
5 employs a recognition algorithm to select one of the arcs A_1 through A_N , and A_{GM} , as the best match, or optionally identifies no match (i.e., if no speech is detected). If A_{GM} is the best arc, the input is rejected as invalid.

With reference now to FIG. 4, the training process will be described. Initially, a main training 207 is performed to derive each utterance, or state
10 model, A_1 through A_N , to be stored in the memory 208, as indicated in step 402. A number of different methods are known for creating the HMM models. In the illustration of FIG. 4, each arc is a left to right, HMM model with no state skips, such that only self loops and single steps transitions are allowed. A brief description of the derivation of such a model is described hereinbelow.
15 Those skilled in the art will recognize that the arcs can be of other known models, and by other known methods.

Initially, features are extracted, in feature extractor 204. It is envisioned that the feature extractor will generate cepstral and delta cepstral coefficients for each frame of an utterance. Those skilled in the art will recognize that there
20 are many ways of calculating cepstral features and for estimating their derivative, and any suitable technique for deriving these coefficients can be used. Frames F_1 through F_N (FIG. 5) are produced during the window, each frame comprising features. Some of the frames represent noise, from which noise energy features are produced by the feature extractor. Other frames
25 represent a portion of the speech signal.

Returning to FIG. 4, in step 604, the processor 108 during training 207, calculates a noise feature for each arc model as indicated in step 604. The noise measurement is made from the feature vectors produced during the start and end of the capture window. In particular, it is desirable to use the average of the
30 feature vectors measured during a start period and an end period of the utterance. For example the first 160 ms, *Savge*, and the last 160 ms, *Eavge*, of the capture window can be used. The capture window is shown in FIG. 5,

including the start period and the end period during which noise feature vectors are stored. The capture window may be 2 seconds long, for example, representing the maximum duration of a word. This capture window can be fixed or variable length, depending on the expected length of the input utterances and the implementation's memory constraints.

The processor 108, having derived the noise feature in step 404, determines whether the device is in hands-free mode in step 404. The device may include a state flag that indicated that the device is in hands-free mode which is activated by the user through a keypad menu or it may include a mechanical connector that actuates a switch when the device 100 is connected to a hands-free kit.

If the device is not in a hands-free mode, the processor calculates during training (which is done independently for each utterance), a noise feature X_{nz} which is the minimum of S_{avg} and E_{avg} (i.e., $\min(S_{avg}, E_{avg})$), as indicated in step 410. For each frame of input speech, an energy value can be computed from its samples. S_{avg} and E_{avg} are averages of these energy values from the indicated frames. The minimum is used for each of the training utterances to update a running noise mean. This noise mean is updated iteratively using the following equation:

$$X_{ref}(k) = ((k-2) * X_{ref}(k-2) + (X_{nz1} + X_{nz2})) / k$$

where $X_{ref}(k)$ is the reference value for the k-th noise feature, X_{nz1} indicates the noise feature found from the minimum of S_{avg} and E_{avg} of the first training utterance and X_{nz2} is the noise feature from the minimum of S_{avg} and E_{avg} of the second training utterance.

The updated noise mean and the number of training utterances used for noise mean updates are recorded in memory 110 as indicated in step 412.

If it was determined in step 406 that the device was in hands-free mode, as indicated in step 408, a hands-free flag HF is set as indicated in step 408. The flag HF is set indicating the presence of

hands-free word models instead of updating the noise model, if training is in hands-free mode.

It is assumed that the training environment will be relatively quiet. This can be enforced through a signal quality check which requires that all training utterances have at least an 18 dB signal to noise ratio. Checks can also be employed to insure that the user does not speak during *Savg* and *Eavg* measurement time.

The general operation of recognition 206 by processor 108 is described generally with respect to FIG. 6. Initially, the noise feature is calculated for the test utterance, which is the input utterance that the system is trying to identify, as indicated in step 602. In the recognition mode, background noise measurement is made from the same initial 160 ms *Savg* and final 160 ms *Eavg* of the utterance window. The noise measurement during recognition is *Xrecog* and is equal to the average of *Savg* and *Eavg*. This value is compared to the reference noise value as calculated in the training mode. A comparison is used to find the ratio of the recognition background noise estimate to the training background noise estimate. Those skilled in the art will recognize that that other relative comparison of these values can be used.

The processor 108 next calculates the word penalty in step 606. The ratio is used to calculate a word entrance penalty. The word entrance penalty controls the strictness of the Out-of-Vocabulary rejection. In general, higher noise environments have a lower strictness value. The word entrance penalty is calculated using a look up table with the noise index ratio being the address for the memory table and the penalty being the output. An advantageous ten penalty distribution as illustrated in FIG. 8 can be used, wherein significantly noisier environments in the recognition mode (ratios 6-9) have a substantially smaller penalty than ratios representing recognition modes closer to the training mode noise reference (ratios 0-4). For example, the curve can be derived as follows:

$$x = X_{ref}(k) / X_{recog}$$

$$f(x) = 1 / (1 + 2^{1.5(x-5)}).$$

Out of range index ratios will default to the minimum word entrance penalty, which is zero. The actual penalty applied may for example be
 5 $-220 * f(x)$, although the actual scalar can be of any value that results in a penalty have a desirable proportion to the scores it is combined with. The use of the non-linear relationship provides a significant improvement of in vocabulary and out of vocabulary recognition by providing a large penalty when noise conditions are good and a small
 10 penalty when noise conditions are bad. Those skilled in the art will recognize that the calculation of the word entrance penalty may be made directly, rather than through the use of a look-up table.

The recognition continues with its main search and parallel garbage models, as indicated in step 608. The goal of the recognition
 15 system is to find the most likely path from node N_1 to Node N_2 in FIG. 3. The nodes N_1 and N_2 are coupled by paths $A_1 - A_N$ representing the Hidden Markov Models for the N word vocabulary, optionally including a garbage model A_{GM} . Additionally A_1^{noise} and A_2^{noise} represent the noise models and are associated with nodes N_1 and N_2 .
 20 The garbage model attempts to capture any non-vocabulary sounds or words in the input utterance. It is a one state zero-valued model used only by the Out-of-Vocabulary rejection algorithm. To prevent it from modeling noise better than the noise model, a penalty is applied to garbage model probability scores for frames classified as noise.

25 The search through the grammar network, as illustrated in FIG. 3, is done by a two level alignment algorithm, such as a Viterbi algorithm. The lowest level of this search finds the best alignment and path score between the frames of the input utterance and the states of a given arc. An example of techniques used to apply frames of an
 30 utterance to states of an individual model are disclosed in copending patent application Docket Number CS10103, entitled METHOD OF TRACEBACK MATRIX STORAGE IN SPEECH

RECOGNITION SYSTEM, filed in the name of Jeffrey Arthur Meunier et al. on the same date as this application, and copending patent application filed on even date herewith, docket number CS10104, entitled METHOD OF SELECTIVELY ASSIGNING A
 5 PENALTY TO A PROBABILITY ASSOCIATED WITH A VOICE RECOGNITION SYSTEM, filed in the name of Daniel Popper, the disclosures of which is incorporated herein by reference thereto. The lower level alignment algorithm generates a score for the best path of the input utterance through the given HMM arc.

10 In addition to the lower level alignment algorithm wherein the scores of each arc, or HMM, are tracked via cumulative probabilities $c_i^n(m)$, which is the cumulative probability of state i of arc A_n at frame m , nodes N_1 and N_2 must also track their own cumulative probabilities. The node cumulative probability $C_j(m)$, is the cumulative
 15 probability of Node N_j at frame m . This probability is calculated much like the cumulative probability of each HMM in that it keeps the highest score to the node. The cumulative probability can be calculated as follows:

$$C_j(m+1) = \text{Max}_{n \in A_j} \{C_{in}^n(m) + \text{Po}_{in}(d_{in})\}$$

20 where A_j is the set of arcs $\{A_1, A_2, \dots, A_N\}$ which terminate at node j , in is the number of states in arc n , d_{in} is the duration of the last state of arc n , and $\text{Po}_1(d_{in})$ is the out of state transition penalty for the last state of arc n . The cumulative probability is the maximum over all arcs that terminate on node N_j of the sum of the last state's cumulative
 25 probability $C_{in}^n(m)$ with its out of state probability $\text{Po}_{in}(d_{in})$.

While tracking the cumulative probabilities for the nodes, the calculation of the cumulative probability for the initial state of each arc, $c_1^n(m)$ must be modified to allow for transitions into its initial state from a node N_j . There is a one time transition penalty assigned to the
 30 transition from the node N_j to the initial state of arc A_n , called a word entrance penalty. It does not apply to the noise model or to the garbage

model, so it acts as a strictness control on the Out of Vocabulary rejection when enabled. The cumulative probability can be seen to be

$$C_i^n(m+1) = o_i^n(f_m) + \max(C_j(m) + W(n), C_i^n(m) + Ps_i(d_1))$$

5 where $W(n) = \begin{cases} g(x) & \text{if } n \in \{A_1, A_2, A_3\} \\ 0 & \text{if } n \in \{A_1^{noise}, A_2^{noise}, A_{GM}\} \end{cases}$
 where $W(n)$ is the word entry penalty, A_{GM} is the garbage arc, A_i^{noise} is the noise arc for node i , $o_i^n(f_m)$ is the observation probability of the feature vector f_m in state i of arc n , and $Ps_i(d_1)$ is the same state transition penalty of state i of arc n . This equation keeps the maximum
 10 of either the same state transition to the transition from the originating node and adds to the observation probability. The information retained at the end of the recognition process is the arc that was traversed to get to node N_2 . This is done by propagation path information along with the cumulative probabilities $C_i^n(m)$ and $C_j^n(m)$.

15 For valid utterances, the word model's best path through the alignment algorithm must produce a better score than the garbage model by a value greater than the word entrance penalty, or the valid utterance will be falsely rejected. For invalid utterances, the garbage model must be greater than the path through each of the eligible word
 20 models such that the utterance is correctly rejected.

 The recognition algorithm uses the entire window of feature vectors collected, which may for example typically be 2 seconds worth of data. Additionally it uses a speech/noise classification bit for each frame to update the one state noise model used in A_1^{noise} and A_2^{noise} of

25 FIG. 3

 In the recognition mode, the processor 108 initializes recognition by setting the noise update flag to 1 and the frame count to zero, as indicated in step 702. The frame count is incremented in step 704. The processor then determines whether the noise flag is set in step
 30 706. If not, the processor proceeds to decision 716. If the flag is set, the processor 108 determines whether the noise model should still enabled in step 708. If not, the noise update flag is set to 0 in step 714. Noise modeling is turned off after a certain number of updates are made

If noise updating should still be performed, the processor determines whether to update the noise model in step 710. If the processor is to update the noise model for the frame, the model is update in step 712. . The noise model A_1^{noise} and A_2^{noise} are computed
5 dynamically by the system through the use of the speech/noise classification bits sent in by the feature extraction algorithm. The details of the decision of whether to update the noise model for the current frame is made by looking at the speech classification made by the feature extraction algorithms. Once a predetermined number of
10 consecutive speech frames are seen for the utterance, no more updates are made. For example, the limit may be 3 frames. The noise model will only be updated for a particular frame if that frame's speech to noise classification indicates that it is a noise frame.

The processor then determines whether the frame count is less
15 than a threshold number of frames in step 716. Probability estimation will not begin until a certain number of frames have been processed. This is to allow the noise model to become somewhat accurate before probabilities based on the noise model are calculated. If the threshold number of frames have not been received, the processor returns to step
20 704 wherein the frame count is incremented by one.

If the frame count exceeds the threshold, the processor 108 calculates cumulative probabilities for the nodes and arcs for the frame in step 718. The probability scores are normalized in step 720. Normalization can be provided by subtracting the largest cumulative
25 probability from all other cumulative probabilities. The cumulative normalization factor is also tracked so that the unnormalized score can be returned at the end of the recognition process.

The processor then determines if the last frame was processed in step 722. If not, the processor returns to step 704 and increments the
30 frame count. Otherwise, the recognition result is output with the normalized score as indicated in step 724.

The noise model is a one state model. The vector mean of this state is $\mu_l^{noise}(m)$ which is a function of m because it is computed dynamically and is updated with a new feature vector f_{m+1} at frame m+1 as follows:

5
$$\mu_l^{noise}(m) = ((M_{noise}(m) * \mu_l) + f_{m+1}) / (M_{noise}(m) + 1)$$

where $M_{noise}(m)$ is the number of noise frames that have been used in the computation of $\mu_l^{noise}(m)$, which can be different that the value of m since not all frames are used in the noise update. Additionally, the update equation is used only for the cepstral elements of the noise model. The delta-cepstral and the delta energy elements are fixed at zero.

Accordingly, it can be seen that an improved system is disclosed providing variable rejection strictness depending upon the background noise levels during training and recognition. The system helps to prevent association of invalid utterances with stored speech models and helps improve the accurate detection of valid utterances.

Although the invention has been described and illustrated in the above description and drawings, it is understood that this description is by way of example only and that numerous changes and modifications can be made by those skilled in the art without departing from the true scope of the invention. Although the present invention finds particular application in portable wireless devices such as cellular radiotelephones, the invention could be applied to any device employing speech recognition, including pagers, electronic organizers, computers, and telephony equipment. The invention should be limited only by the following claims.

CLAIMS

1. A method of operating a voice recognition system,
comprising the steps of:
 - 5 generating a variable rejection strictness as a function of at least one background noise level measured during training and noise signal measurements made during an input utterance made during recognition mode of operation; and
 - deriving a word entrance penalty as a function of the variable
10 rejection strictness.
2. The method as defined in claim 1, wherein the step of
generating a variable rejection strictness includes the step of measuring
noise during at least a portion of the training utterance for a model.
15
3. The method as defined in claim 1, further including the
step of selectively updating the noise features from the training
utterances.
- 20 4. The method as defined in claim 1, further including the
step of storing noise statistics during training with a model so that they
are available to the recognition algorithm.
5. The method as defined in claim 3, wherein noise
25 statistics are not updated when training in a hands-free mode.
6. The method as defined in claim 3, further including the
step of generating a signal to noise ratio, and wherein training is
prohibited if the signal to noise ratio is below a predetermined level.
30
7. The method as defined in claim 1, wherein during
recognition, if no noise statistics are available for an utterance, the

recognition algorithm defaults to the minimum strictness requirement when applying the alignment algorithm to that utterance.

8. The method as defined in claim 1, wherein during
5 recognition, the input noise energy feature is compared to the reference noise statistics and a noise ratio is computed.

9. The method as defined in claim 8, wherein the strictness
of the out of vocabulary rejection algorithm is then selected based
10 upon the noise ratio.

10. The method as defined in claim 1, wherein the
confidence measurement of the best path is implemented using a zero
mean one state garbage model in parallel with the voice tag models.
15



Application No: GB 0003269.8
Claims searched: 1 to 10

Examiner: John Donaldson
Date of search: 14 March 2000

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
UK CI (Ed.R): G4R(REX, RRL)
Int CI (Ed.7): G10L 15/00, 15/08, 15/14, 15/20
Other: Online:WPI, EPODOC, JAPIO

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	GB 2188764 A (NRDC), see abstract	-

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.